The Cielo Petascale Production Capability System

Manuel Vigil, HPC-DO

LANL and SNL have formed a partnership, the New Mexico ACES, to acquire and deploy a production capability system for the Department of Energy NNSA ASC program to support the national SSP.

Following the contract award to Cray in March 2010, the New Mexico Alliance for Computing at Extreme Scales (ACES)/Cray partnership has achieved significant progress under a very aggressive schedule (see Fig. 1) in preparation for tri-lab simulations work in 2011. This is the first time a multi-lab partnership has been involved in deploying an Advanced Simulation and Computing (ASC) program capability machine. This partnership has worked well in helping meet all contractual, project, and program milestones.

Effective management and stewardship of the nuclear weapons stockpile into the future requires the ability to accurately assess the behavior of weapons to ensure robust and reliable performance while maintaining the testing moratorium. These accurate assessments drive the requirements for predictive capability in weapons science, including a fine-scale numerical resolution and advanced models for physics and material behavior. The National Nuclear Security Administration's (NNSA) ability to assess the systems in the stockpile must rely on science-based prediction. These capabilities are required for supporting the Stockpile Stewardship Program (SSP) certification and assessments, ensuring that the nation's nuclear stockpile is safe, reliable, and secure.

Fig. 1. Cielo 2009-2011 Schedule.

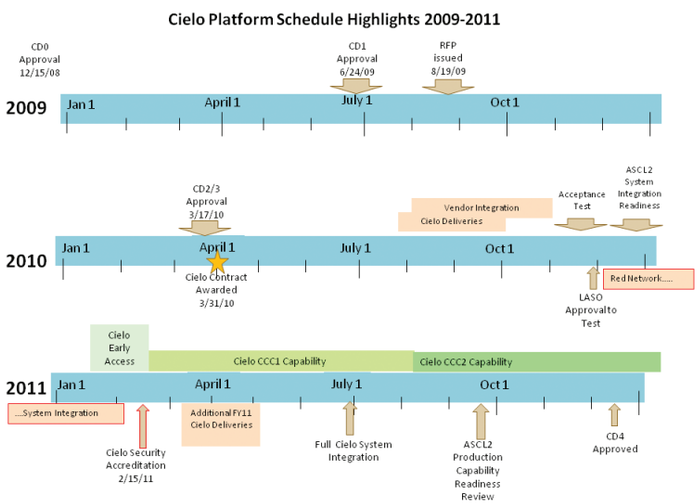

The Cielo platform is targeted to be the signature production classified capability platform resource for running integrated weapons simulations for the tri-labs (LLNL, LANL, and SNL) in the 2011–2015 time frame. The Cielo platform provides a replacement computing resource for existing simulation codes, as well as providing a larger resource for ever-increasing capability computing requirements. Cielo will be sited at LANL, but will operate under a national user facility paradigm by the ASC program and will be available to the tri-lab community.

The major objectives for this system are: (1) to provide a petascale production capability system capable of running a single application across the entire machine, (2) to serve as the ASC flagship capability machine for running the ASC Capability Computing Campaign (CCC) simulations, (3) to provide an application performance emphasis for running the current suite of ASC integrated codes, running capability class jobs, and provide a greater than 6-fold performance improvement over the previous capability machine (Purple at LLNL), and (4) to proceed with an emphasis on easy migration of existing integrated weapons codes. In developing the requirements for this system, ACES personnel solicited input from system and application users at LLNL, LANL, and SNL. System and applications technical experts also participated in the technical evaluation of proposals for the selection of the system.

In March, 2010, Cray, Inc. was selected to deliver the Cielo platform for ACES. The deliveries for Cielo were set up to provide a ~1.03 Petaflop/s system in FY10 and additional deliveries in FY11 for a total system peak capability of ~1.37 Petaflop/s. The initial system consists of 6704 nodes and uses a 3D Torus topology and Cray's Gemini interconnect. A node is composed of two 8-core AMD Magny Cours processors with 16 GB of memory, resulting in a total of 107,262 cores and 160 TB of memory. Cielo will also be the first large Cray system to use the Panasas file system for storage. Several types of nodes are supported, including compute, service, and visualization nodes. Visualization nodes were

For more information contact Manuel Vigil at mbv@lanl.gov.

Fig. 2. The first row of the Cielo system now installed at LANL.

acquired with twice the memory of other compute nodes. The initial FY10 system of 72 racks is already installed at the Nicholas C. Metropolis Center for Modeling and Simulation at LANL.

The successful deployment and stabilization of a large computing system requires a strong partnership between the vendor providing the system and the ACES integration team. Since the contract award in March 2010, there has been a lot of effort to build, deliver, test, and accept the Cielo system in anticipation of early application runs in January 2011 (see Fig. 1). ACES is responsible for the overall integration and operations of the platform after final system acceptance.

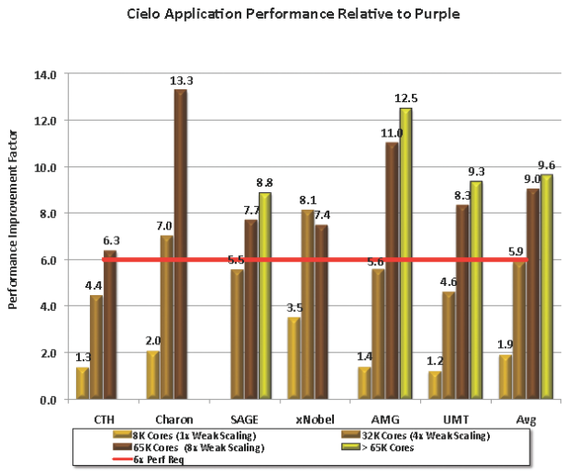

Fig. 3. Six selected ASC applications were tested on Cielo. The final average performance improvement factor was 9.6x; performance was almost ten times faster on Cielo than on Purple.

The activities associated with the acquisition consist of contract award, vendor design and system build, vendor project risk plans, pre-ship and post-ship testing, system delivery, installation at the LANL site, and acceptance testing.

Other efforts include project management, system integration, facilities upgrades, site preparation, file system infrastructure deployment, on-site analyst support, stabilization testing, applications porting and testing, performance tuning and testing, and preparations for production operations of the system, including providing maintenance after the initial warranty period.

Under a very aggressive schedule the ACES/Cray system integration team has made significant progress in 2010 in getting the system ready for applications use. All the original project schedules have been met thus far. A few high-level accomplishments include:

- 72 Cray cabinets built, delivered, and installed at LANL
- System integration in the classified network, including scaling testing
- Panasas system delivered, installed, and tested in the unclassified network
- Completion of official acceptance test
- Transitioning the system to the LANL classified network for integration

In particular, the application performance improvements exhibited by the selected tri-lab application codes have exceeded the contractual requirements (Fig. 3). The early applications work has also demonstrated that most codes tested thus far have been able to port to Cielo with minimal difficulty.

As of January 2011, the system is being further tested in the classified network in preparation for some early application simulations scheduled for mid-January through mid-February of 2011.

Each Laboratory (LLNL, LANL, SNL) has selected an early application for demonstrating the use of Cielo and to help with system stabilization. After expected security accreditation in mid-February, the Cielo Capability Computing Campaign 1 (CCC-1) will begin. This period will be used to run more tri-lab simulations, to help tri-lab application code teams develop or port applications to Cielo, and to allow for application scaling work on those applications in preparation for submitting CCC-2 proposals.

Much work remains to be done before Cielo is in production mode, but the work thus far indicates that Cielo will be an excellent addition to the ASC computing capabilities for supporting the stockpile stewardship work for the tri-lab user community.

A system upgrade to add the additional 24 Cray cabinets to the Cielo system is scheduled for the second quarter of 2011.

Funding Acknowledgments
DOE NNSA Advanced Simulation and Computing (ASC) program; in support of Stockpile Stewardship Program (SSP)